

Searches for Axion-Like Particles with NGC1275: Observation of Spectral Modulations

Marcus Berg,¹ Joseph P. Conlon,² Francesca Day,² Nicholas Jennings,² Sven Krippendorf,² Andrew J. Powell,² Markus Rummel²

¹Department of Physics, Karlstad University, 651 88 Karlstad, Sweden

²Rudolf Peierls Centre for Theoretical Physics, 1 Keble Road, Oxford, OX1 3NP, UK

E-mail: marcus.berg@kau.se, joseph.conlon@physics.ox.ac.uk,
francesca.day@physics.ox.ac.uk, nicholas.jennings@physics.ox.ac.uk,
sven.krippendorf@physics.ox.ac.uk, andrew.powell2@physics.ox.ac.uk,
markus.rummel@physics.ox.ac.uk

Abstract. Axion-like particles (ALPs) can induce localised $\mathcal{O}(10\%)$ oscillatory modulations in the spectra of photon sources passing through astrophysical magnetic fields. Ultra-deep *Chandra* observations of the Perseus cluster contain over 5×10^5 counts from the central AGN, NGC1275, and represent a dataset of extraordinary quality for ALP searches. We use this dataset to search for X-ray spectral irregularities from the AGN. The absence of irregularities at the $\mathcal{O}(30\%)$ level allows us to place leading constraints on the ALP-photon mixing parameter $g_{a\gamma\gamma} \lesssim 1.5 - 5.9 \times 10^{-12} \text{GeV}^{-1}$ for $m_a \lesssim 10^{-12} \text{eV}$, depending on assumptions on the magnetic field realisation along the line of sight. At $\mathcal{O}(10\%)$ level two modulations are present at high statistical significance, an excess in the 2–2.2 keV region and a deficit at 3.4–3.5 keV. We are unable to account for these through conventional instrumental or astrophysical processes and, interpreted as a signal, they would correspond to an ALP-photon coupling in the range $g_{a\gamma\gamma} \sim 1 - 5 \times 10^{-12} \text{GeV}^{-1}$.

Contents

1	Introduction	1
2	AGNs and ALP physics	3
3	The Observations	6
3.1	Chandra	6
3.2	XMM-Newton	7
4	Chandra Analysis	8
4.1	Pile-Up: General Comments	8
4.2	Full Spectra	10
4.3	Accounting for Pile-up	13
5	XMM-Newton	20
5.1	Analysis	20
5.2	Pile-up in XMM-Newton	23
6	Bounds	24
7	Potential Signal	26
7.1	Pile-Up	26
7.2	Effective Area Miscalibration	28
7.3	Gain Miscalibration	29
7.4	Astrophysical Explanations for 2–2.2 keV excess	30
7.5	Astrophysical Explanations for 3.4–3.5 keV deficit	31
7.6	ALP interpretation of features	33
8	Conclusions	33
A	Ultraclean Spectra	35

1 Introduction

The Perseus galaxy cluster (A426) is the brightest X-ray cluster in the sky. It is a cool-core cluster at redshift $z = 0.0176$, centered around the Seyfert galaxy NGC1275 and its Active Galactic Nucleus (AGN). Due to its proximity and brightness, the Perseus cluster has been

a standard target for all X-ray satellites. The X-ray spectrum and emission specifically from the AGN are described in [1–4].

Axion-like particles (ALPs) are a well-motivated extension of the Standard Model, and arise generically in string compactifications (for example see [5–7]). As the potential and interactions of ALPs are protected by shift symmetries, they can naturally have extremely small masses or even be massless.¹ If it exists, an ALP a interacts with the Standard Model via a coupling to electromagnetism,

$$\frac{a}{M} \mathbf{E} \cdot \mathbf{B} , \quad (1.1)$$

where $M \equiv g_{a\gamma\gamma}^{-1}$ parametrises the strength of the interaction and \mathbf{E} and \mathbf{B} are the electric and magnetic fields. A general review of ALPs and their physics can be found in [8]. We concern ourselves in this paper with the case $m_a \lesssim 10^{-12} \text{eV}$, for which non-observation of gamma ray photons coincident with the SN1987A neutrino burst constrains $M \gtrsim 2 \cdot 10^{11} \text{GeV}$ [9–11].

If ALPs exist, then the interaction of equation (1.1) causes ALPs and photons to interconvert in the presence of a background magnetic field $\langle B \rangle$ [12, 13]. Starting with a pure photon spectrum, this $\gamma \leftrightarrow a$ interconversion results in modulations in the spectrum of arriving photons.

Galaxy clusters are particularly efficient photon-ALP converters [14, 15], and for the electron densities and magnetic fields present within galaxy clusters, it is a result that at X-ray energies the $\gamma \leftrightarrow a$ conversion probability is both energy-dependent and quasi-sinusoidal [15–18]. Compared to the source spectrum, the spectrum of arriving photons then has oscillatory modulations imprinted on it. By searching for such modulations, we can place constraints on the coupling parameter M .

For this purpose, quasars or AGNs that are either behind or embedded in galaxy clusters provide attractive sources. The original photon spectrum is reasonably well described by an absorbed power-law, and all photons arise from a single sightline passing through the cluster. As bright sources, AGNs can also provide the large number of counts necessary for statistical leverage in searching for oscillatory modulations of the photon spectrum.

This method is indeed already largely described and was used in [16] (see also [19, 20], and [21] for a recent analysis of NGC1275 in gamma rays using this approach). However, [16] only applied these ideas to the study of the AGN at the centre of the Hydra A galaxy cluster (redshift $z = 0.052$), for which 200ks of *Chandra* observation time exist. But in this case, the combination of the intrinsic AGN brightness and the redshift of $z = 0.052$ results in only a few thousand counts in total, limiting the ability to produce bounds.

In contrast, the AGN at the centre of NGC1275 is both exceedingly bright and the subject of enormous observational time. NGC1275 has been observed for 1 Ms by *Chandra* ACIS-S, with observations taken in 2002 and 2004, and also for a further 0.5 Ms by *Chandra* ACIS-I in 2009. Taken together, these generate over half a million X-ray counts from the central AGN – a dataset of extraordinary quality for searching for spectral irregularities. Furthermore, there are also 180 ks of observation time with *XMM-Newton* taken in 2001 and 2006. This dataset is not as rich as the *Chandra* dataset, but allows us to cross check our analysis with a different instrument. In this paper, we use these datasets to search for, and constrain, axion-like particles.

¹The shift symmetry also implies that, even if massless, ALPs are not constrained by searches for fifth forces or modifications of general relativity.

The paper is organised as follows. Section 2 provides further details on the physics of ALPs and the attractiveness of bright quasars or AGNs for searching for ALPs. Section 3 describes the *Chandra* and *XMM-Newton* observations that we have used and their properties. Section 4 describes the *Chandra* analysis, the effects of pile-up and strategies taken to mitigate pile-up. Section 5 describes the analysis of *XMM-Newton* data. In Section 6 we describe bounds on the ALP-photon coupling and in Section 7 we discuss a potential ALP signal and candidate astrophysical or instrumental origins. In Section 8 we conclude.

2 AGNs and ALP physics

If axion-like particles exist, they interconvert with photons in a background magnetic field $\langle B \rangle \neq 0$ [12, 13, 22]. This occurs because the magnetic field generates, via equation (1.1), a 2-particle interaction between the ALP a and the photon γ , resulting in a mixing of the mass eigenstates. In a mathematically identical fashion to neutrino oscillations, photons and ALPs then have a finite probability of inter-conversion as they pass through the magnetic field. This is a quantum-mechanical effect, and results in an evolution of the quantum state

$$|\psi_{init}\rangle = |\gamma(E)\rangle \longrightarrow |\psi_{final}\rangle = \alpha|\gamma(E)\rangle + \beta|a(E)\rangle ,$$

where $|\alpha|^2 + |\beta|^2 = 1$, and $\gamma(E)$ [$a(E)$] denotes a photon [ALP] with energy E . The conversion probability $|\beta|^2$ depends on the free electron density (which sets the effective photon mass), the magnetic field (which sets the strength of the mixing) and the magnetic field coherence length (which determines the region over which mixing applies).

The calculational details of this are standard. Following the original papers [12, 13, 22], the physics of axion-like particles passing through magnetic fields has been described in many works. An incomplete list of articles studying aspects of photon-ALP interconversion in astrophysical magnetic fields includes [11, 14–18, 23–33].

In this respect, galaxy clusters at X-ray energies sit at a sweet spot for photon-ALP physics. This is due to two key results. First, galaxy clusters are *particularly efficient* environments for photon-ALP interconversion. The electron densities are relatively low. Clusters have magnetic fields that are not significantly smaller than in galaxies, but in which the \mathbf{B} -field extends over megaparsec scales, far greater than the tens of kiloparsecs applicable for galactic magnetic fields. The magnetic field coherence lengths in clusters are also larger than in galaxies, comfortably reaching tens of kiloparsecs. For massless ALPs, this feature singles out galaxy clusters as providing the most suitable environment in the universe for ALP-photon interconversion.²

For convenience we will restrict to massless ALPs in this paper. As the efficiency of ALP-photon conversion depends on $|m_a^2 - \omega_{pl}^2|^{-2}$, where ω_{pl} is the plasma frequency in galaxy clusters, in practice all masses $m_a \lesssim 10^{-12}\text{eV}$ are equivalent for our purposes. For $m_a \sim 10^{-12}\text{eV}$, resonant conversion can occur, allowing weaker ALP-photon couplings to be probed, but we will not consider that case explicitly in this paper.

The second key result is that, for the electron densities and magnetic field structures present within galaxy clusters, the photon-ALP conversion probability is energy-dependent,

²Although they appear appealing, magnetars and related objects do not provide efficient environments for ALP-photon conversion [13].




Figure 1. Left—The photon survival probability along a line of sight modelled on that from NGC1275 to us, for a randomly generated magnetic field. A central magnetic field of $B_0 = 15\mu G$ was assumed, with a radial scaling of $\langle B(r) \rangle \sim n_e(r)^{0.7}$. There were 100 domains, with lengths drawn randomly between 3.5 and 10 kpc, and the total propagation length being 620 kpc. The ALP-photon coupling is $g_{a\gamma\gamma} = 1.5 \times 10^{-12} \text{GeV}^{-1}$ (roughly a factor of three beyond the current upper limit $g_{a\gamma\gamma} < 5 \times 10^{-12} \text{GeV}^{-1}$ from SN1987A). One should not take too seriously the particular details of the magnetic field model used, as there is no observational data to constrain the precise range of magnetic field coherence lengths in the Perseus cluster, but the quasi-sinusoidal structure arises generically. Right—The photon survival probability for the same magnetic field convolved with a Gaussian with FWHM of 150eV.

with a quasi-sinusoidal oscillatory structure at X-ray energies. This provides distinctive spectral features to search for. We illustrate this in Figure 1, where we plot a typical photon survival probability as a function of energy, along a single line of sight modelled on that from NGC1275 to us.

The precise form of the survival probability is not predictable. It depends on the actual magnetic field structure along the line of sight, and so differs for each line of sight. Faraday rotation measures can give statistical information about the strength and extent of magnetic fields, as for the Coma cluster in [34], but the actual magnetic field along any one line of sight is unknown. However the form shown in Figure 1 – a quasi-sinusoidal structure with a period that increases with energy – is generic, and arises for any reasonable choice of central magnetic field value or range of coherence lengths. The inefficiency of conversion at energies $E \lesssim 0.2 \text{keV}$ is also generic, implying that effects of photon-ALP conversion are not visible in the optical (and below) range.

Figure 1 (right) also shows the same survival probability convolved with a Gaussian with full width at half maximum (FWHM) of 150eV, representing the approximate energy resolution of the CCD detectors present on *Chandra* and *XMM-Newton* satellites (the precise figure of 150eV is taken from the in-orbit performance of the ACIS-I detectors on *Chandra*, see table 6.4 of the *Chandra* proposer’s guide³). While at lowest energies the oscillations are too rapid to be resolved by CCD detectors, and would require micro-calorimeters such as those that were present on *Hitomi*⁴, in general it is fortuitous that the scales of the oscillations match those of the X-ray telescopes extensively used to observe galaxy clusters.

If ALPs exist, then for photons arriving from a single location, this conversion imprints

³<http://cxc.harvard.edu/proposer/POG/html/chap6.html>

⁴The energy resolution from *Hitomi* is expected to be $\sim 5 \text{ eV}$.

a particular quasi-sinusoidal modulation on the actual photon spectrum. There is also an overall reduction in luminosity, but this can be absorbed into the overall normalisation of the spectrum. For unpolarised light the $\gamma \rightarrow a$ conversion probability cannot exceed fifty per cent, and in the limit of strong coupling saturates at an average value of $\langle P(\gamma \rightarrow a) \rangle = 1/3$ (for example, see [19]). It therefore follows that, expressed as a ratio of data to model, the maximal allowed range of ALP-induced modulations is approximately $\pm 30\%$.⁵

There are three effects that can wash out these modulations. The first is the finite energy resolution of the telescope; as shown in Figure 1, this removes any structure present at the lowest energies. The second is when emission arises from an extended source, involving many different lines of sight. In this case the peaks and troughs from different lines of sight undergo destructive interference, reducing any signal. The third is insufficient photon statistics, when localised oscillations become indistinguishable from Poisson fluctuations.

Bright point-like sources either behind or embedded in a galaxy cluster are particularly attractive for searching for ALP-induced modulations. The galaxy cluster provides a good environment for ALP-photon conversion; the bright point source ensures there are many photons, all passing along the same line of sight.

These factors make quasar or AGN spectra attractive for searching for ALPs. Emission from an active galactic nucleus (AGN) arises from matter accreting onto the central black hole. As evidenced by the rapid time variability of AGN luminosities, the physical region sourcing the X-ray AGN emission is tiny – of order a few Schwarzschild radii of the central black hole. As cluster magnetic fields are ordered on kiloparsec scales, this implies that for all practical purposes every photon arising from the AGN has experienced an identical magnetic field structure during its passage to us.

To first approximation, at X-ray energies an AGN spectrum can be described as an absorbed power-law. The effect of ALPs is then to imprint a quasi-sinusoidal modulation on this power law, of relative amplitude at most $\mathcal{O}(10\%)$ and with a modulation period of order a few hundred eV. As the fractional Poisson error on N counts is $\frac{1}{\sqrt{N}}$, and CCD detectors such as those on *Chandra* and *XMM-Newton* have intrinsic energy resolutions of around $\mathcal{O}(100\text{eV})$, it therefore requires large numbers of counts to be able to distinguish any ALP-induced modulations from normal statistical fluctuations.

All the above facts make the AGN of the Seyfert galaxy NGC1275 an excellent candidate for searching for ALP-photon interconversion. NGC1275 is the central galaxy of the Perseus cluster, which as a cool core cluster should have a high central magnetic field (estimated as $25 \mu\text{G}$ in [35]) – implying the sightline from NGC1275 to us should be efficient at ALP-photon conversion.

The nucleus is intrinsically bright and unobscured, with a spectrum that is well characterised by a power-law and narrow Fe $K\alpha$ line, absorbed by the galactic n_H column density [1]. Furthermore, there is enormous *Chandra* observation time on NGC1275, encompassing 1.5 Ms in total. This results in over half a million photon counts originating from the central AGN, although for the on-axis observations quite a number are contaminated by pile-up. This is a level two orders of magnitude larger than the study in [16] involving Hydra A.

⁵We re-emphasise here that photon-ALP conversion involves quantum oscillations between states rather than absorption. Therefore for passage from $A \rightarrow C$ the survival probability $P(A \rightarrow C)$ does *not* equal $P(A \rightarrow B) \times P(B \rightarrow C)$.




Figure 2. NGC1275 in three types of observation, from left to right: centre of the chip in 2004 (*Chandra* ObsID 4952), midway between the edge and centre of the chip in 2009 (ObsID 11714), and the edge of the chip in 2009 (ObsID 11713). The colour coding is adjusted to account for the different observation times such that each colour corresponds to the same count-rate across images.

3 The Observations

3.1 Chandra

The deep Chandra observations involving NGC1275 can be divided into three main groups. The first involves 200ks of ACIS-S observations taken in 2002 together with 800ks of ACIS-S observations taken in 2004.⁶ In these observations NGC1275 is close to the aimpoint. The second group involves 300ks of ACIS-I observations carried out in 2009, where NGC1275 is approximately midway between the edge of the chips and the aimpoint. The third group involves 200ks of ACIS-I observations also taken in 2009, in which NGC1275 is close to the edge of one of the chips, around 8 arcminutes from the aimpoint. Finally, there are also some brief pre-2002 observations that we do not include.

The relevance of this classification is that the point spread function of the telescope grows off-axis. In the first group, the photons from the AGN suffer little dispersion and are highly concentrated on a few pixels. In the third group, the arriving photons are scattered over many pixels, whereas the second group is intermediate.

The consequence of this is that, despite having the shortest observation time, it is the third group (the 200ks in which NGC1275 is at the edge of the detector) that provides the cleanest data set. In the first case, the superb optics of *Chandra* work against it; almost all photons are concentrated onto a few pixels, and these central pixels are highly contaminated by pile-up. In this case a clean spectrum can only be obtained by extracting from the wings of the point spread function – which however reduces the photon count. In the last case however, being highly off-axis causes sufficient degradation of the optics that the arriving photons are scattered over many pixels, resulting in greatly reduced pile-up. The second grouping is intermediary in quality between these two. This is illustrated in Figure 2, which shows images of NGC1275 for each of the different observation types.

Another relevant factor is that the brightness of the AGN varies substantially with time. As described in [4], the NGC1275 AGN was brightest from 1970 to 1990, before

⁶Note for particle theorists: the ACIS instrument has two main modes, ACIS-S and ACIS-I. One ACIS-S chip leads to an 8 arcminute by 8 arcminute field of view, while ACIS-I will result in a 16 arcminute by 16 arcminute field of view.

Obs ID	Exposure [ks]	Year	Instrument	Location of NGC1275
3209	95.77	2002	ACIS-S	Central
3404	5.31	2002	ACIS-S	Central
4289	95.41	2002	ACIS-S	Central
4946	23.66	2004	ACIS-S	Central
4947	29.79	2004	ACIS-S	Central
6139	56.43	2004	ACIS-S	Central
6145	85	2004	ACIS-S	Central
4948	118.61	2004	ACIS-S	Central
4949	29.38	2004	ACIS-S	Central
6146	47.13	2004	ACIS-S	Central
4950	96.92	2004	ACIS-S	Central
4951	96.12	2004	ACIS-S	Central
4952	164.24	2004	ACIS-S	Central
4953	30.08	2004	ACIS-S	Central
11715	73.36	2009	ACIS-I	Midway
11716	39.64	2009	ACIS-I	Midway
12037	84.63	2009	ACIS-I	Midway
11714	91.99	2009	ACIS-I	Midway
11713	112.24	2009	ACIS-I	Edge
12025	17.93	2009	ACIS-I	Edge
12033	18.89	2009	ACIS-I	Edge
12036	47.92	2009	ACIS-I	Edge

Table 1. The Chandra observations used in this paper. The last column shows the location of NGC1275 in the respective observation relative to the focal point.

rapidly declining by an order of magnitude until around 2000. Since then it appears to have brightened significantly over the decade from 2003 onwards, although it has not yet returned to the luminosities it had pre-1990.

The combination of pile-up and the intrinsic brightening of the AGN implies that, despite the shorter observation times, each of the stacked 200ks and 300ks ACIS-I edge and midway observations from 2009 has more counts than the 1Ms of ACIS-S observations taken in 2002 and 2004. Furthermore, when we consider cleaned spectra that exclude regions of high pile-up, the total counts in the ACIS-I edge spectra is more than in the ACIS-I midway and ACIS-S observations put together. These factors make the ACIS-I edge observations the optimal for analysis purposes, even though they involve the shortest observational time.

3.2 XMM-Newton

There are two observations of NGC1275 with significant exposure time, see Table 2. The first one, 0085110101, was taken in 2001 when the AGN had its lowest emissivity in obser-

Obs ID	Exposure [ks]	Year	Instrument	Location of NGC1275
0085110101	53.08	2001	EPIC	Central
0305780101	123.3	2006	EPIC	Central

Table 2. The *XMM-Newton* observations used in this paper. The last column shows the location of NGC1275 in the respective observation relative to the focal point.

vational history. The second observation, 0305780101, was taken when the emissivity was still relatively low but nevertheless almost twice as bright as in 2001. NGC1275 is on-axis in both observations.

The 2001 observation was taken in full frame mode, while the 2006 observation was taken in extended full frame mode. This affects the frame time of the pn camera, which is 73.4 ms in full frame mode and 199.1 ms in extended full frame mode. For MOS, the frame time is 2.6 s. The pixel size of pn is 4.1 arcseconds and 1.1 arcseconds for MOS. This means that pn in extended full frame mode is more susceptible to pile-up than MOS. In general, pile-up is an issue for all *XMM-Newton* observations of NGC1275 as we will discuss in the *XMM-Newton* analysis Section 5.

A significant difference between *XMM-Newton* and *Chandra* is the angular resolution. For both MOS and pn the radius of the disk containing 50 % of the photons collected in the focal plane (Half Energy Width) is around 8.5 arcseconds at 1.5 keV, while for *Chandra* it is much smaller at ~ 0.5 arcseconds. As the central region of the Perseus cluster is also intrinsically bright, this makes it harder to separate AGN and cluster emission for the *XMM-Newton* observations.

The effective area at 1 keV is 922 cm² for MOS and 1227 cm² for pn compared to 340 cm² for *Chandra*. This allows MOS and pn combined to collect roughly 7 times more photons in a given observation time than *Chandra* (although this also increases the amount of pile-up). However, there is roughly 10 times more exposure time in our *Chandra* dataset and the significantly better angular resolution of *Chandra* allows a much better contrast to be attained between the AGN and the cluster emission. This leads to much better statistics in our *Chandra* dataset in terms of the total number of counts after background subtraction.

4 Chandra Analysis

We use CIAO 4.7 [36], Sherpa [37] and HEASOFT 6.17 for the *Chandra* data analysis.⁷ After the data is reprocessed using CALDB 4.6.9, it is cleaned from time periods that are polluted by flares using the program chips. We find that only Obs ID 4950 is affected by flares and the cleaning reduces the observation time slightly from 96.12 ks to 89.23 ks for this observation.

4.1 Pile-Up: General Comments

As the brightness of the AGN makes pile-up a significant feature of all observations, we will make some general comments on it here.

⁷The update to CIAO 4.8 affects data taken in Continuous Clocking mode, which does not apply to these observations.

The energy recorded on the individual ACIS pixels (each approximately 0.5 arcseconds square) is read out approximately every three seconds (one frame time) in these observations. Based on groupings of 3×3 pixels, events are graded. Pile-up refers to the arrival of more than one photon in this grouping within the same readout frame. This can lead to the energy of the two (or more) incident photons being summed, and either treated incorrectly as a single photon event of higher energy, or assigned a bad grade (grade migration). For an on-axis bright source (as in the ACIS-S observations of NGC1275), the level of pile-up can be high, and the resulting spectrum contains events with two, three (and more) photons. As pile-up is a statistical feature of the number of arriving photons, some level of pile-up is inevitable in any observation. The question is always whether the magnitude of pile-up is sufficient to corrupt the science analysis being undertaken.

In terms of the measured photon distribution, the general effect of pile-up is to cause a hardening of the spectrum: two or more lower-energy photons are misidentified as a single higher-energy photon. This implies that for a fit of a single power-law to a photon distribution, as pile-up increases the best-fit power-law index will decrease. In a spectrum contaminated by pile-up, this makes it harder to determine the correct original power-law index.

What about searches for and constraints on ALPs? As we have seen in Section 2, the distinct signal of ALPs is a quasi-sinusoidal modulation in the spectrum – a local excess or deficit in the photon count rate compared to the nearby continuum. For sufficiently small levels of pile-up, localised modulations will remain localised modulations, as an overall global continuum redistribution of photons is unable to create or remove localised spikes (or dips) relative to the continuum.⁸ However, for sufficiently heavy pile-up, the majority of photons will be redistributed and such local features will be lost.

While pile-up is always a contaminant on the spectrum, what this implies is that a search for localised spectral irregularities is much more robust against pile-up than, for example, a measurement of the overall power-law spectral index of a source. While some movement of counts from low to high energies will reduce the number of low-energy photons compared to high energy photons, it will be less likely to affect the presence or absence of sharp localised features.

This robustness is more applicable at low energies. The effective area of the *Chandra* telescopes starts falling rapidly above around 5 keV.⁹ As at higher energies a power-law distribution also produces intrinsically fewer photons, it only requires a small amount of pile-up of lower-energy photons into the $E > 5$ keV region to cause a significant distortion of the spectrum there.

In contrast, at lower energies the effective area is larger and there are far more photons, so small amounts of pile-up will not affect any spectral features. For observations towards Perseus, there is an additional benefit: the high galactic absorbing column density ($n_H = 1.5 \times 10^{21} \text{ cm}^{-2}$) removes the lowest-energy photons, resulting in an effective minimal value for a piled-up energy $E_1 + E_2$ (this will be relevant when we discuss in detail below a feature at 2–2.2 keV).

There is an important caveat to this which requires careful treatment. In the presence of

⁸In the case of a strong low-energy emission line, this is not true, as its pile-up may result in spikes at integer multiples of the original line.

⁹cf. cxc.harvard.edu/proposer/POG/html/ACIS.html

rapid variations in the effective area, a failure to account for pile-up can result in significant spectral distortion. This is because the distribution of pile-up photons around $E_1 + E_2$ does not respect the behaviour of the effective area at $E_1 + E_2$, but instead only the effective areas near E_1 and E_2 . If the effective area is rapidly varying near $E_1 + E_2$, and pile-up photons make up a large fraction of the total photon count at $E_1 + E_2$, then a fitted spectrum will introduce unphysical features at $E_1 + E_2$. This will be particularly relevant for the *XMM-Newton* spectra.

The level of pile-up differs greatly between the three observational groups. In the ACIS-S observations, the central pixels are very heavily affected by pile-up. For a full extraction of the ACIS-S spectrum (with no exclusion of a central region) over 5% of photon counts have energies in the 10-15 keV range. As the effective area of the telescope is zero at these energies, these counts all arise from piled-up events in which several photons have arrived in the same readout frame. In contrast, pile-up is far more moderate for the ACIS-I edge observations (for which a similar complete extraction results in only 0.1 % of counts lying in the 10-15 keV range).

There are two basic methods we can employ to reduce the effects of pile-up on the analysis. The first is to reduce the amount of pile-up by using an annular extraction region and extracting the spectrum only from the wings of the point spread function. The point spread function of the telescope causes the arriving photons to be spread out on the detector, with the degree of spread increasing as one moves progressively off-axis. Furthermore, the point spread function is mildly energy dependent: photons with higher energy are spread out further than photons with smaller energy. While fewer photons arrive in the wings of the point spread function, those that do suffer less from pile-up than those arriving on the central pixels. The second is to model the pile-up, and the distortion it induces on the spectrum, explicitly. We shall utilise both approaches in Sections 4.3.1 and 4.3.2 below.

The overall balance here is between enhanced photon statistics – but with more pile-up and so with worse data quality – and fewer photons but better data. In particle physics language, this is the trade-off between efficiency and purity. We however start with analyses involving complete extractions of the spectra, with no exclusions of high pile-up regions.

4.2 Full Spectra

We first extract complete spectra for NGC1275 for the ACIS-I edge observations (11713, 12025, 12033, 12036). For all the following analyses, spectra and responses were created using `specextract` for each observation, and then stacked using `combine_spectra`. An ellipse around NGC1275 of radii 11.6 and 7.2 arcseconds was used for the extraction region. The background is taken from an elliptical annulus around NGC1275, with the outer radii being 19 and 13 arcseconds and the inner radii 13.3 and 9.3 arcseconds (for the two short observations 12025 and 12033 this region goes beyond the edge of the chip, and a rectangular box was used instead for the background).

The resulting stacked spectrum contains around 266000 counts, reducing to 230000 after background subtraction, giving a ratio of 6.5:1 for the AGN against the cluster emission. After background subtraction, the fraction of counts in the 7–10 keV band and 10–15 keV bands were 1.3 % and 0.09 % respectively (we use these percentiles as a loose measure of the extent of pile-up contamination). The resulting spectrum was binned to ensure a minimum of 2000




Figure 3. The complete stacked spectrum of the ACIS-I edge observations, involving 229000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 2.1 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 1.77$. σ refers to the standard deviation from the model expectation for a Poissonian count rate.

counts a bin,¹⁰ and fitted between 0.8 and 5 keV with an absorbed power law `xswabs × powlaw1d`,

$$AE^{-\gamma} \times e^{-n_H \sigma(E)}. \quad (4.1)$$

Here A denotes the normalisation of the power-law, γ the power-law index, and n_H the effective Hydrogen column density.

The resulting fit is shown in Figure 3, together with the fractional ratio of data to model. The best-fit value of n_H is $2.1 \times 10^{21} \text{cm}^{-2}$ and the power-law index is $\gamma = 1.82$. While the absorbed power-law is a reasonable characterisation of the data, there are two large localised residuals: one positive between 2–2.2 keV and one negative around 3.4–3.6 keV. There is an upward trend at 5 keV. As the effective area of *Chandra* begins to fall off rapidly here, and there are also intrinsically fewer photons expected, pile-up plays a proportionately more important role. This rising trend continues beyond 5 keV and we attribute this to the effects of pile-up.

We next plot the complete spectra for the ACIS-I observations in which NGC1275 is midway on the chip (11714, 11715, 11716, 12037). As for each observation NGC1275 is in different locations relative to the aimpoint, the optical distortions differ for each case and customised extraction regions were used. For 11714, this was a circle of radius 3.5 arcseconds. For 11715, this was an ellipse of radii 4.6 and 6.4 arcseconds. For 11716, an ellipse of radii 4.4 and 5.7 arcseconds was used, while for 12037 an ellipse of radii 4.1 and 6.4 arcseconds

¹⁰In general we bin so that there are approximately one hundred bins in total. If there are too few counts per bin, then the fit is insensitive to localised modulations as can be produced by ALPs, as the goodness-of-fit is insensitive to the sign of the residuals.




Figure 4. The complete stacked spectrum of the ACIS-I central observations, involving 242000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.0 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 1.30$.

was used. In all cases the background was taken from a circular annulus around NGC1275 of radii 8.5 and 20 arcseconds.

The resulting spectrum contains around 259000 counts, reducing to 242000 after background subtraction, giving an AGN to cluster ratio of 14.2:1. After background subtraction, the fraction of counts in the 7-10 keV band and 10-15 keV bands are 3.4 % and 0.6 % respectively. The resulting spectrum was binned to ensure a minimum of 2400 counts per bin, and fitted between 0.8 and 5 keV with an absorbed power law.

The resulting fit and the ratio of data to model are shown in Figure 4. The best-fit parameters are $n_H = 1.0 \times 10^{21} \text{cm}^{-2}$ and spectral index $\gamma = 1.30$. There is a large overall modulation in the data to model ratio that arises from pile-up contamination, but superimposed on it there is a clearly visible large localised positive residual just below 2.2 keV. The data to model ratio for this residual is entirely consistent with the same feature in the ACIS-I edge data.

Finally, we extract spectra for the ACIS-S observations in which NGC1275 is centrally located. A circle of radius 3.4 arcseconds is used for the extraction region, and the background is taken from an annulus of inner and outer radii 5 and 12 arcseconds. The resulting stacked spectrum contains around 233000 counts, reducing to 183000 after background subtraction, giving a ratio of 3.7:1 of AGN to cluster emission. After background subtraction, the fraction of counts in the 7-10 keV band and 10-15 keV bands are 10.4 % and 6.7% respectively (evidence of the very heavy pile-up present for these on-axis observations). The resulting spectrum was binned to ensure a minimum of 1800 counts per bin, and fitted between 1.5 and 5 keV with an absorbed power law.




Figure 5. The complete stacked spectrum of the ACIS-S observations, involving 183000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.9 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 0.66$.

The resulting fit is shown in Figure 5 together with the data to model ratio. The best-fit parameters are $n_H = 1.9 \times 10^{21} \text{cm}^{-2}$ and an extremely hard, pile-up induced, spectral index of $\gamma = 0.66$. As with the ACIS-I midway observations, there is nonetheless a sharp localised excess in the 2–2.2 keV region superimposed on the pile-up.

4.3 Accounting for Pile-up

The above spectra show interesting features. Despite the very different levels of pile-up, there is evidence at high statistical significance for a $\sim 10\%$ excess in the region 2–2.2 keV, and also strong evidence in the ACIS-I edge observations for a similar deficit in the region 3.4–3.5 keV.

However it is certainly true that all these spectra are, to varying degrees, affected by pile-up. To mitigate this, in this section we describe two ways of reducing the effects of pile-up. The first method involves excluding a central region of highest pile-up from the analysis. While this reduces the photon counts – and so the statistical significance of any features in the spectrum – it also allows produces purer spectra of higher quality. The second method is to keep the full spectrum, but model the pile-up explicitly.

4.3.1 Cleaned Spectra

Here we describe the creation and analysis of spectra with a central region of high pile-up excluded.

We analyse the images and count rates of each observation using the image software `ds9`. The precise method of central exclusion is somewhat arbitrary; the method used to produce




Figure 6. The stacked spectrum of the ACIS-I edge observations, with central pixels excluded according to the prescription in the text. There are 153000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 2.2 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 1.85$.

the spectra was as follows. We set the ds9 binning such that the ds9 pixels are essentially the same size as the physical *Chandra* pixels (half an arc-second across). We then manually exclude all ds9 pixels which are neighbours (either side-by-side or diagonally) to any pixel with total counts greater than 1% of the overall exposure time. Using the `pileup_map` tool, this results in almost all retained pixels having a pile-up fraction lower than 5%.

We first show in Figure 6 the ACIS-I edge observations with a central exclusion according to the above method. As these observations are relatively clean to begin with, the fractional change in the number of counts is relatively small. There are now 187000 counts before background subtraction and 153000 counts afterwards. The fraction of counts in the 7–10 keV and 10–15 keV regions are 1.2% and 0.04 %. The best-fit spectral index is $\gamma = 1.85$ with $n_H = 2.2 \times 10^{21} \text{cm}^{-2}$. While the statistical significance of the features at 2–2.2 keV and 3.4–3.5 keV reduces (consistent with the reduced photon counts), the magnitudes of the data to model fluctuations remain the same.

We now perform a similar cleaning of the ACIS-I observations with the source midway on the chip. As the central pixels are heavily piled-up, in this case the cleaning procedures significantly increases the quality of the fit. This is at the cost of a significant reduction in photon statistics: there are now only 88000 counts before background subtraction, and 74000 after background subtraction (with 0.1 % in the 10–15 keV band). We group counts so that there are 700 counts per bin, leading to a spectral index of $\gamma = 1.64$ with $n_H = 1.3 \times 10^{21} \text{cm}^{-2}$. With a Q-value of 0.18 and a reduced χ^2 of 1.14 for 86 degrees of freedoms, this is a good fit to the data. The cleaned spectrum still shows a clear excess in the 2.1–2.2 keV region at a ratio consistent with other observations. While not at notable significance, there is also a mild local deficit at 3.4–3.5 keV. In both cases, the ratio of these features in terms of




Figure 7. The stacked spectrum of the ACIS-I midway observations, with central pixels excluded according to the prescription in the text. There are 74000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.3 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 1.64$.

data to model is entirely consistent with the high-statistics results from the full ACIS-I edge observations.

We now repeat this procedure for the ACIS-S observations. In this case we reduce pile-up by removing a central square of 16 pixels from the extraction region. The resulting cleaned spectrum has 117000 counts before background subtraction and 74000 counts after background subtraction. After background subtraction there are now 0.1% counts in the 10–15 keV band and 1.1% of counts in the 7–10 keV band, indicating that this spectrum is now substantially cleaner. We group counts so that there are at least 700 counts per bin. In this case, an absorbed power-law is not sufficient for a good fit and we supplement this by a soft thermal component using *xsapec* (the presence of a thermal component for NGC1275 was also reported in [2]). The presence of a soft thermal component substantially improves the fit (taking the Q value from 10^{-18} to 10^{-2}). The resulting fit has $n_H = 1.7 \times 10^{21} \text{cm}^{-2}$ and a power-law index of $\gamma = 1.85$. The thermal component has a temperature $T = 0.85 \text{keV}$. At this temperature, the amplitude and abundance of the thermal component are largely degenerate in the fit. Fixing the abundance at solar abundance, the relative amplitude of thermal component to the power-law is 0.15. Again, the cleaned ACIS-S spectrum displays a clear positive residual centred around 2.1 keV and a clear negative residual at 3.4 keV, with amplitudes that are consistent with those found in the ACIS-I edge spectrum. The residuals also show an oscillatory structure, although given the error bars one should not over-interpret this.

Why is a soft thermal component necessary for a good fit in the ACIS-S observations but not in the ACIS-I observations? There are two main reasons. First, the effective area of ACIS-S has more support at low energies than for ACIS-I. The net result is that for




Figure 8. The stacked spectrum of the ACIS-S observations, with a central 16-pixel square excluded. The fit is to an absorbed power law and soft thermal component.

ACIS-S the peak of the observed photon spectrum is around 1 keV, whereas for the ACIS-I observations the peak is around 1.6 keV. This implies that the ACIS-S observations have a proportionally higher fraction of low energy photons, and so are more sensitive to the presence of a soft thermal component. Secondly, the ACIS-S and ACIS-I observations were taken at different times - the former in 2002 and 2004 and the latter in 2009. Over this period the intrinsic brightness of the AGN has increased significantly. Assuming the absolute strength of the soft thermal component to have remained the same, its relative significance would be much greater in the ACIS-S observations.

We have checked that one can add a soft thermal component to the ACIS-I spectra, of the same absolute magnitude as for the ACIS-S spectra, without affecting the quality of the fit. As for the ACIS-I observations the thermal component neither significantly improves or worsens the fit, and does not affect the above results, we have not included it explicitly for the ACIS-I data.

4.3.2 Modelling Pile-up

The previous subsection addressed pile-up by removing central, high-pile-up regions from the extraction region. Here we take a complementary approach, aiming to model the pile-up directly. While this modelling does not provide a complete picture of pile-up, it ameliorates its effect. A full treatment would require customised tools going beyond the scope of this paper (the pile-up tool `jdpileup` [38] provided with Sherpa is optimised for on-axis sources, while NGC1275 is off-axis in all ACIS-I observations).

We first describe our pile-up modelling for the ACIS-S observations, where NGC1275 is on-axis. This results in a high degree of pile-up – the central pixels have large numbers of

counts above 10 keV, all of which arise from multi-photon pile-up events.

We use the `jdpileup` model as described in [38]. As this model assumes the source is on-axis, it is directly appropriate here. The model assumes Poisson statistics to calculate the probability of different numbers of photons hitting an event-detection cell (a 3×3 pixel region in Chandra) within the read-out time adjusted to these respective observations. For the ACIS-S and ACIS-I midway observations the read-out time was 3.1 seconds, and for the ACIS-I edge observations it was 3.2 seconds. The model then convolves this probability with the probability of such events being assigned a ‘good’ grade, and the conversion from photon energy to pulse height by the detector. Obviously this is a non-linear process: the pile-up in any particular bin depends on the energy spectrum for all energies below the bin. The parameters of the model therefore need to be determined together with those for the spectrum under consideration. This can lead to degeneracies in parameter space, in particular for a simple power law.

As per [38] the two parameters of the pile-up model that we allow to vary are α and f , where α^{p-1} is the probability that p piled photons will be assigned a ‘good’ grade, and f is the fraction of events to which pile-up will be applied. The other parameter that will prove important is n , the number of regions to which `jdpileup` will be applied independently. For a point source this should be set to 1, as was done for the ACIS-S observations. For an extended source it should roughly correspond to the number of 3×3 pixel islands in the region. The reason for this is that `jdpileup` assumes spatial uniformity across the extraction region. For the ACIS-I midway observations, where the AGN is smeared out across several pixel islands, the value of n proved difficult to determine for large extraction regions, with the value corresponding to the best fit being unrelated to the number of pixel islands. We therefore constrained ourselves to a smaller central region, with little variation in count rate between pixels, to give us more control over the pile-up model.

It is also worth noting that the `jdpileup` model is set to zero for energies less than 0.5 keV and energies greater than 10 keV. As there are many events above the 10 keV range for the ACIS-S observations, we only model the spectrum up to 10 keV, as extending the fit beyond that would result in the model parameters being sent to unphysical values.

For the ACIS-S pile-up analysis, we used a circular region around NGC1275 of radius 3.5 arcseconds, with the background taken from an annular region of inner radius 4 arcseconds and outer radius 7 arcseconds. Counts were grouped to 1500 per bin. The spectrum was modelled using an absorbed power law and thermal emission with temperature $T = 0.85\text{keV}$, and fitted with the `jdpileup` model for energy values between 1 and 10 keV. The data was fit using the `moncar` Monte-Carlo method, and run several times to ensure the global minimum had been found.

The resulting fit and the ratio of data to model are shown in Figure 9. The best fit parameters involve $n_H = 2.6 \times 10^{21}\text{cm}^{-2}$ and a spectral index $\gamma = 1.81$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.660$ and $f = 0.943$ respectively. We can clearly see that this model provides a good description of the data all the way up to 10 keV, and produces physically sensible values for n_H and γ , despite the model estimating a pile-up fraction of 82%. The reduced statistic of the fit is 1.75, with a Q-value of 10^{-5} . As before, we again see a clear excess around 2.0–2.2 keV, and a low-significance dip at 3.4–3.5 keV. There is also a narrow local deficit around 4.6 keV at around 3-sigma significance. The 6.3 keV iron line (6.4 keV in the cluster rest frame) is obscured by the high level of pile-up in that energy




Figure 9. The ACIS-S observations, involving 177000 counts after background subtraction. The fit is to an absorbed power law with a thermal component of $T = 0.85$ keV, and pile-up is modelled with `jdpileup`. The ratio of data to model is shown at the bottom of this figure.

region.

While we re-emphasise that this will not represent a perfect account of pile-up, it does capture the relevant physics, producing a sensible fit with physical parameters – and with the same features as seen in the far cleaner ACIS-I edge spectra.

For the ACIS-I midway observations we considered a central 6×6 pixel extraction region (this is almost exactly the complement of the clean ACIS-I midway spectrum used in the previous subsection). The resulting spectrum contains around 136000 counts, reducing to 134000 after background subtraction, giving a very high AGN to cluster contrast of 67:1. After background subtraction, the fraction of counts in the 7–10 keV band and 10–15 keV bands are 5.3% and 1.2% respectively. Counts were grouped to 1000 per bin. The spectrum was modelled using an absorbed power law and fitted with the `jdpileup` model, this time for energy values between 1 and 9 keV, to ensure no counts with energy greater than 10 keV were included in the final bin. The `jdpileup` parameter n was set to 4, the number of 3×3 pixel islands in the extraction region.

The resulting fit and ratio of data to model are shown in Figure 10. The best fit parameters involve $n_H = 2.5 \times 10^{21} \text{cm}^{-2}$ and a spectral index $\gamma = 1.93$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.324$ and $f = 0.975$ respectively, and the estimated pile-up fraction was 35%. Once again the excess around 2.0–2.2 keV is apparent, along with a small deficit at 3.4–3.5 keV. The (unmodelled) 6.3 keV iron line is also visible as an excess in the data. Again, the fit is good up to 9 keV.

We finally consider the case of the ACIS-I edge observations, where pile-up is relatively weak. Although a reasonable fit can be made with no pile-up modelling, as argued above,




Figure 10. The ACIS-I midway observations from a central 6×6 pixel region, involving 134000 counts after background subtraction. The fit is to an absorbed power law and pile-up is modelled with `jdpileup`.

we now aim at modelling the pile-up also in these observations for completeness. There is a rule of thumb¹¹ that 0.007 counts/second per 3×3 pixel island is about 1% pile-up, and 0.07 counts/second is about 10%. We will assume that the ACIS-I edge observations with mild cleaning have pile-up fractions of 10% or lower.

For the ACIS-I edge observations, we used as source an elliptical annulus of radii 3 and 5 arcseconds, with the 12 brightest pixels removed by a contour. The background was extracted from an elliptical annulus of outer radii 18 and 23 arcseconds, with inner radii 8 and 12 arcseconds. The resulting spectrum contains around 128000 counts, reducing to 111000 after background subtraction, giving an AGN to cluster ratio of 6:1. After background subtraction, the fraction of counts in the 7–10 keV band and 10–15 keV bands are 1.1% and 0.1% respectively. We grouped counts to 1000 per bin. Here there are fewer counts out at higher energies than in the previous two sets of observations, so we fit only out to 5 keV. The `jdpileup` parameter n was set to 16, the number of 3×3 pixel islands in the extraction region. This should only be thought of as a rough estimate, as the `jdpileup` model is stated to be accurate for on-axis point sources, and the edge observations are also not spatially uniform in terms of counts. The fit is to an absorbed power law, and results in $n_H = 2.7 \times 10^{21} \text{cm}^{-2}$ and a spectral index of $\gamma = 1.89$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.52$ and $f = 0.87$ respectively. The pile-up fraction was 3%, consistent with estimates and very mild as expected. The excess around 2.0–2.2 keV is again apparent, as well as the small deficit at 3.4–3.5 keV.

Finally, we also considered *Chandra*’s grating observations of NGC1275 (ObsIDs 333

¹¹cxc.harvard.edu/csc/memos/files/Davis-pileup.pdf




Figure 11. The ACIS-I edge observations, involving 110000 counts after background subtraction. The fit is to an absorbed power law and pile-up is modelled with `jdpileup`.

and 428 with HETG) as a means to reduce pile-up, but this reduced counts to the point where it did not seem worthwhile.

5 XMM-Newton

5.1 Analysis

We use SAS version 15.0.0. The 2001 and 2006 datasets are reprocessed with up to date calibrations via `cifbuild`. To remove flares we apply the standard filters of counts/s < 0.35 for MOS and counts/s < 0.4 for pn. There is a significant flare towards the end of the 2006 observational period and various shorter flares during the 2001 observation. Removing the polluted time intervals results in reducing the effective exposure time from 53 ks to 49 ks for MOS 2001, 25 ks to 7.4 ks for pn 2001, 123 ks to 117 ks for MOS 2006 and 76 ks to 49 ks for pn 2006.

For the extraction regions we first choose circle of radii 13.8 arcseconds (MOS1 2001), 17.5" (MOS2 2001), 14.3" (MOS1 2006), 17.5" (MOS2 2006), 15.5" (pn 2001), and 19" (pn 2006). We then check if pile-up is an issue using the SAS tool `epatplot` and find that pile-up is in general present in all the observations. Since the AGN was less bright by a factor ~ 2 in 2001, pile-up is slightly less of an issue for these observations. On the one hand, the smaller pixel size of 1.1 arcseconds of the MOS cameras makes it less susceptible for pile-up than the 4.1 arcseconds pixel size of the pn camera. On the other hand, pn's time resolution of 73.4 ms in full frame mode is advantageous with respect to pile-up compared to the 2.6 s time resolution of MOS. However, in extended full frame mode pn's 199.1 ms time

resolution make it less advantageous with respect to pile-up compared to MOS. Note that after flare removal there are only 7.4 ks of pn data in full frame mode while there are 49 ks of more piled up data in extended full frame mode. As the MOS data is expected to be more sensitive, we do not present the pn data from either 2001 or 2006 and instead focus on the MOS data only.

In order to make a compromise between avoiding the most heavily piled-up regions and not too small signal over background ratio we choose the following inner annuli radii 2.5 arcseconds (MOS1 2001), 3" (MOS2 2001), 7.5" (MOS1 2006), 7.5" (MOS2 2006). For the background regions we choose annuli of 13.8-22 arcseconds (MOS1 2001), 17.5-22.5" (MOS2 2001), 14.3-22.5" (MOS1 2006), 17.5-22.5" (MOS2 2006). The worse angular resolution of *XMM-Newton* compared to *Chandra* implies that the discrimination of the AGN emission against the cluster is relatively poor (for the 2006 data, the AGN:cluster data ratio is 1:3, whereas for *Chandra* spectra this ratio is not worse than 1.7:1).

After applying `eveselect` to generate the spectra we use `rmfgen` and `arfgen` to generate the response files. We use `epicspeccombine` to combine the MOS1 and MOS2 spectra of the 2001 and 2006 observations, respectively. We also use the `ftools` routines `mathpha`, `addrmf` and `addarf` (`epicspeccombine` does not produce an `arf` file) and find that the results are comparable within statistical errors to using `epicspeccombine`. In order to avoid combining different systematic errors we do not combine the 2001 and 2006 data.

The spectra are further analysed using the spectral fitting programs `Xspec` [39] and `Sherpa` [37], requiring a minimum count-rate per bin of about 1.5% of the overall count rate. The energy interval is 0.5–7.5 keV.

For NGC1275, we use a spectral model of a powerlaw as in our *Chandra* analysis

$$A(E) = K E^{-\gamma}, \quad (5.1)$$

where K is the normalisation in photons $\text{keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$ and γ is the photon index. To model a soft excess of the AGN, we add an `apec` model whose temperature we expect to be around 1 keV. The abundance parameter of this model is set to unity.¹² We also include the well known iron line at 6.4 keV in the cluster rest frame in our spectral fit.

To model photoelectric absorption, the source spectrum is multiplied by the `zwabs` model

$$M(E) = e^{-n_H \sigma(E)}, \quad (5.2)$$

where n_H is the hydrogen column density and σ is the photoelectric cross-section [40]. Photoelectric absorption is mostly relevant below approximately 1 keV if $n_H \gtrsim 10^{21} \text{ atoms cm}^{-2}$ (The galactic column density is $n_H = 1.5 \cdot 10^{21} \text{ atoms cm}^{-2}$.)

The fitted spectra are shown in Figures 12 and 13. We do not observe any significant excess at 2–2.2 keV or a deficit at 3.5 keV, although there is a mild dip in the 2006 MOS data at 3.5 keV, see Figure 13. However, this is not a contradiction with having these features real and present in the *Chandra* data. First of all, the statistics for *XMM-Newton* are more limited: the two spectra, combined, contain a total of ~ 110000 counts for the MOS observations after background subtraction. Secondly, the data is much noisier: the overall

¹²Note that around 1 keV this `apec` spectrum is dominated by atomic lines. Hence, there is a degeneracy between the abundance parameter and the normalisation of the `apec` model.

	MOS 2001	MOS 2006
Exposure [ks]	49	118
Counts	~ 47000	~ 63000
γ	1.65 ± 0.03	1.65 ± 0.03
T_{soft}	0.84 ± 0.06	0.80 ± 0.06
$n_H [10^{22} \text{ atoms cm}^{-2}]$	0.13 ± 0.03	0.13 ± 0.03
χ^2 / dof	154 / 131	205 / 178
q-value	0.08	0.08

Table 3. Fit results for the *XMM-Newton* MOS observations from 2001 and 2006. These results are in agreement with the fit parameters found in [1] for the 2001 MOS data of NGC1275.

spectrum is background dominated, making it harder to ensure sensitivity to small fractional residuals in the AGN spectrum. Finally, pile-up can significantly alter the spectrum of the *XMM-Newton* observations around 2 – 2.4 keV, as we now discuss in Section 5.2.




Figure 12. MOS 2001 spectral fit. The best fit parameters are listed in Table 3.




Figure 13. MOS 2006 spectral fit. The best fit parameters are listed in Table 3.

5.2 Pile-up in XMM-Newton

We describe more precisely here why, in terms of spectral features around 2–2.2 keV, the effects of pile-up are much more severe for *XMM-Newton* than for *Chandra*. The key plot for this is Figure 14, which shows the relative behaviour of pile-up for *Chandra* ACIS-I and *XMM-Newton* MOS cameras. This plot shows, for both instruments, the expected distribution of pile-up events compared to the distribution of single-photon events for observations of NGC1275. The basic assumptions underlying this plot are that pile-up is dominated by two-photon events and that the overall spectrum provides a relatively accurate picture of single-photon events. The pile-up distribution then comes, essentially, from adding the spectrum to itself.

We suppose that a fixed overall pile-up fraction (for definiteness 10%) applies in both cases. What will be the impact around 2 keV?

For *Chandra*, the peak in the pile-up spectrum is around 3 keV. Figure 14 implies that the proportion of piled-up photons around 2–2.2 keV is lower than the proportion of good photons around 2–2.2 keV. This means that the fractional contribution of pile-up to the spectrum in the 2–2.2 keV region is less than the overall pile-up fraction. In contrast, for *XMM-Newton*, the peak in the pile-up spectrum is around 2.4 keV. Above 2 keV, the fractional contribution of pile-up is larger than the overall pile-up fraction. The proportion of piled-up photons in this region is then higher than the overall pile-up fraction. This becomes

most severe around 2.4 keV for *XMM-Newton*, where the pile-up spectrum peaks. Here, a 10% overall pile-up fraction would result in a local additional contribution of $\sim 30\%$ due to pile-up.

This is where the structure of the MOS effective area (cf. [41]) is also crucial. The MOS effective area has edges around 1.8 and 2.3 keV, with the 2–2.2 keV region as a local plateau above these. To leading order, the distribution of single photon counts in this region should respect these dips and plateau in the effective area – and this is the behaviour that a fitting program will attempt to enforce. However pile-up events do not respect the effective area structure at the energy at which they are registered. When a significant number of pile-up events arrives across this region, it distorts the resulting fits. In particular, it generates artificial excesses in regions of low effective area and artificial deficits in regions of high effective area.

As for *XMM-Newton* all observations of NGC1275 are on-axis, pile-up is always a significant contaminant for the central regions where a significant contrast can be achieved against the cluster emission. Fits to this central region produce deficits in the 2 - 2.2 keV and excesses above that. This behaviour is clearest in the 2006 pn observations, which has the largest degree of pile-up, but it is also visible for the MOS cameras.¹³ Consistent with a pile-up origin, this effect reduces with cleaner data samples.

We also note that for *XMM-Newton* observations of NGC1275, pile-up is in general harder to deal with than for *Chandra* and there are no really clean data samples. For central extraction regions, pile-up is a significant contaminant. For annular extraction regions, the large point spread function of *XMM-Newton* makes it hard to obtain a strong contrast of the AGN emission compared to the cluster. Furthermore, even for annular extraction regions pile-up is not eliminated as a contaminant due to the intrinsic brightness of the central region of the Perseus cluster, and the larger pixel sizes on *XMM-Newton*,

6 Bounds

Before considering the potential signals in the data, we want to use the above results to constrain ALPs and the ALP-photon mixing parameter.

It is clear from the spectra in the previous sections that an absorbed power-law is a good fit to the Perseus data up to residuals of around 10%. Using this, we can say that on passing from NGC1275 through the Perseus cluster and towards us, $\langle P(\gamma \rightarrow a) \rangle \lesssim 20\%$. It follows that ALP-photon couplings large enough to generate the saturated limit of $\langle P(\gamma \rightarrow a) \rangle = 1/3$ are excluded, as these would produce larger residuals in the data than are actually observed.

To obtain approximate bounds on $g_{a\gamma\gamma}$, we compare two models for the flux $F(E)$ observed from NGC1275:

- Model 0: An absorbed power law $F_0(E) = AE^{-\gamma} \times e^{-n_H \sigma(E)}$, as described in Equation (4.1).

¹³This happens even when the actual data points shows a rising trend in counts in the 2 - 2.2 keV region. For example, inspection of the data for pattern 0 photons for a central extraction for the 2006 MOS observations shows a clear apparent excess at 2.1 - 2.2 keV, but the overall fit reports a deficit across this whole region.

- Model 1: An absorbed power law multiplied by the photon survival probability assuming the existence of ALPs with coupling $g_{a\gamma\gamma}$. In this case the predicted flux also depends upon the magnetic field \mathbf{B} along the line of sight. We have $F_1(E, \mathbf{B}) = AE^{-\gamma} \times e^{-n_H \sigma(E)} \times P_{\gamma \rightarrow \gamma}(E(1+z), \mathbf{B}, g_{a\gamma\gamma})$.

Although we have (limited) empirical estimates of the strength of the magnetic field in Perseus, the exact structure is unknown. In practice, we randomly generate many instances of the field from a given power spectrum. The parameter most relevant to ALP-photon conversion is the central magnetic field strength B_0 , estimated as $25 \mu\text{G}$ in [35]. We assume that B decreases with radius as $B \propto n_e^{0.7}$. The electron density n_e has the radial distribution found in [1],

$$n_e(r) = \frac{3.9 \times 10^{-2}}{[1 + (\frac{r}{80 \text{ kpc}})^2]^{1.8}} + \frac{4.05 \times 10^{-3}}{[1 + (\frac{r}{280 \text{ kpc}})^2]^{0.87}} \text{ cm}^{-3}.$$

We simulate each field realisation with 100 domains. The length l of each domain is between 3.5 and 10 kpc, randomly drawn from a Pareto distribution with minimum length 3.5 kpc and power 2.8. We therefore have:

$$P(l > x) = \begin{cases} 0 & \text{for } x > 10 \text{ kpc} , \\ (\frac{3.5 \text{ kpc}}{x})^{2.8} & \text{for } 3.5 \text{ kpc} < x < 10 \text{ kpc} , \\ 1 & \text{for } x < 3.5 \text{ kpc} . \end{cases} \quad (6.1)$$

The coherence length and power spectrum of the magnetic field in the centre of Perseus is not observationally determined. Instead, these parameters are motivated by those found for the cool core cluster A2199 [42], taking a conservative value for the magnetic field radial scaling. The magnetic field and electron density are constant in each domain, with $B(r)$ and $n_e(r)$ evaluated at the centre of the domain and the direction of \mathbf{B} chosen at random.

We compute 95% confidence limits on $g_{a\gamma\gamma}$ by generating fake data from Model 1 and assessing how well it is fit by Model 0 i.e. how well the oscillations due to ALP-photon conversion can hide in the Poisson noise. We use the clean ACIS-I edge observations for this analysis. We fit the spectrum between 1 and 4 keV (a region unaffected by pile-up) and bin such that there are 1000 counts in each energy bin. We use Sherpa's Levenberg-Marquardt fitting method with Poisson errors derived from the value of the data in each bin. Our procedure to determine whether ALPs with coupling $g_{a\gamma\gamma}$ are excluded at the 95% confidence level is as follows:

1. Fit Model 0 to the data and find the corresponding reduced χ^2 , χ_{data}^2 .
2. Randomly generate 50 different magnetic field realisations \mathbf{B}_i for the line of sight to NGC1275.
3. For each \mathbf{B}_i , compute $P_{\gamma \rightarrow \gamma}(E, \mathbf{B}_i, g_{a\gamma\gamma})$ by numerically propagating photons at different energies through \mathbf{B}_i , as described for example in [17]. We take 300 photon energies equally spaced between 1 and 4 keV.
4. For each \mathbf{B}_i , generate 10 fake data sets from Model 1, using Sherpa's fake pha method.

5. Fit Model 0 to each of the 500 fake data sets and find the corresponding reduced χ^2 , χ_i^2 for each.
6. If fewer than 5% of the χ_i^2 are lower than χ_{data}^2 , $g_{a\gamma\gamma}$ is excluded at the 95% confidence level.

We scan over $g_{a\gamma\gamma}$ in steps of $10^{-13} \text{ GeV}^{-1}$. For the $g_{a\gamma\gamma}$ value excluded, we also check that the three values above it in our grid are also excluded. For the magnetic field parameters described above, we find $g_{a\gamma\gamma} \lesssim 1.5 \times 10^{-12} \text{ GeV}^{-1}$. If we consider a more pessimistic scenario with $B_0 = 15 \mu\text{G}$ and a minimum coherence length of 0.7 kpc, we instead find $g_{a\gamma\gamma} \lesssim 3.8 \times 10^{-12} \text{ GeV}^{-1}$. If we take an even more pessimistic scenario in which the central field is $B_0 = 10 \mu\text{G}$ and the minimum coherence length is 0.7 kpc, the bound increases further to $g_{a\gamma\gamma} \lesssim 5.9 \times 10^{-12} \text{ GeV}^{-1}$.

7 Potential Signal

The *Chandra* data shows two main features departing from a power law: an excess at 2–2.2 keV, present in all *Chandra* observations and with overwhelming statistical significance, and a deficit 3.4–3.5 keV that is not as strong but is still present at almost five (local) sigma in the ACIS-I edge data. These features ensure that an absorbed power law (plus thermal component) is not a good fit to the data. We first want to consider possible instrumental or astrophysical explanation for these features.

7.1 Pile-Up

We first consider the possibility that pile-up is responsible for the presence of these features. As the excess around 2–2.2 keV arises at the same location as a significant dip in the *Chandra* effective area, the role of pile-up deserves serious scrutiny. The presence of an edge in the effective area at $E_1 + E_2$ is irrelevant to the distribution of pile-up events where two photons of energies E_1 and E_2 are registered as an event of energy $E_1 + E_2$.

However, there are many arguments that these features do not arise from pile-up.

1. Most simply, for the NGC1275 spectrum it is actually rather difficult (particularly with the ACIS-I detectors) to create piled-up events with energies in the 2–2.2 keV region. This is due to a combination of the large absorbing galactic column density towards Perseus and the effective area of the detectors. The combined effect of these is that there are few soft photons in the spectrum; for the ACIS-I spectra, the peak in the count rate is around 1.5 keV, which implies that it is hard to create pile-up photons with energies in the 2–2.2 keV region.

We illustrate this in Figure 14. This shows the probability distribution of piled-up 2-photon events in the spectrum compared to the probability distribution of the full spectrum. We see that only a very small fraction of pile-up events occur in the 2–2.2 keV region. Furthermore, we see that the overall number of pile-up events is rising with energy up to around 3 keV. As illustrated in the model plots in Section 4, a dip followed by a rise in counts with energy in this region is also the expected feature of the model, and also observed in the background spectra. However the 2–2.2 keV




Figure 14. The distribution of pile-up 2-photon events (blue line) compared to the distribution of single photons, for the *Chandra* ACIS-I observations (left), and also for the *XMM-Newton* MOS camera (right). We derive these by assuming that the full spectrum is a good approximation to the spectrum of single-photon events, which holds for small or moderate degrees of pile-up. The 2-photon distribution is then taken as the sum of the 1-photon distributions. The main point to notice is that in *Chandra* the relative proportion of pile-up photons near 2 keV is much smaller than for *XMM-Newton*, and the pile-up peak is higher, at around 3 keV.

excess is manifest as a continual fall in the photon count with energy – which is not the behaviour produced by pile-up.

2. Secondly, the spectra studied involve very different amounts of pile-up. If these features were due to pile-up, then we would expect their strength to depend on the amount of pile-up, and to disappear for clean spectra with minimal pile-up.¹⁴ However, this is not the case. For example, the 2–2.2 keV feature is instead present as a consistent 10–15 % excess in the data to model ratio for all spectra, ranging from the full ACIS-S spectra to the clean ACIS-I edge spectra, with the fraction of photons that are piled up changing by an order of magnitude (in Appendix A we provide a further spectrum illustrating this fact, where the data is extracted from deep in the wings of the point spread function, reducing pile-up to the 1–2% level). It is hard to see how this could be compatible with an origin from pile-up.
3. Thirdly, the features remain even when we model the pile-up. Although the pile-up model will not capture every aspect of the data, it does produce a good fit even for the heavily piled-up ACIS-S data and clearly correctly captures many of the consequences of pile-up. However, even when explicitly modelling the pile-up the features still remain at the same level.

¹⁴This is what happens with the *XMM-Newton* data, where there is a strong spurious local excess around 2.4 keV in the most highly piled up pn data – which reduces and then disappears when using data that has smaller amounts of pile-up.

4. Although we do not present plots, we have also examined the data keeping only grade 0 photons. This is the cleanest possible *Chandra* photon selection and greatly reduces the proportion of photons that are piled up. A restriction to purely grade 0 photons is not supported by the standard *Chandra* calibrations and so cannot be used for detailed fits; however, visual inspection of this data shows that the excess at 2–2.2 keV and deficit at 3.4–3.5 keV still remain when only grade 0 photons are used.

As all techniques used to reduce or model pile-up do not seem to affect either the presence or strength of these features, we conclude that it is unlikely that they are caused by pile-up.

7.2 Effective Area Miscalibration

Another possibility could be that there is a local miscalibration in the effective area. Certainly, the 2–2.2 keV feature does arise precisely at the location of an effective area edge, where the effective area drops rapidly. A 10 percent miscalibration in the effective area in the 2–2.2 keV regime would be able to reproduce the excess we observe there.

However, such an instrumental effect does not seem plausible given the following checks.

1. We have observed the same excess at a consistent ratio of 10% across the three separate groups of observations. These involve two separate instruments (ACIS-S and ACIS-I), with NGC1275 located at very different locations on the chips. In one case the AGN is located on the optical axis, in another case it is midway between the optical axis and the edge of the chip, and in a third case it is on the very edge of the chip. By itself, it would be hard to exclude the possibility of a localised contaminant near the edge of the telescope affecting the ACIS-I edge observations. However, we see the same features for all different locations of NGC1275 – including the on-axis ACIS-S observations for which the effective area is extremely well calibrated.
2. We have also analysed data from the main part of the Perseus cluster, using the same observations that we have used in our study of NGC1275. Thermal bremsstrahlung provides (as expected) a good fit to the data, and no unexpected residuals are seen in either the 2–2.2 keV or 3.4–3.5 keV region. In particular, in the 2–2.2 keV the data has the expected behaviour of a dip in the number of photon counts, followed by a rise.

We have also utilised the deep *Chandra* observations (500ks) of the Coma cluster taken in 2012 to perform a similar test. As the Coma cluster is relatively isothermal, thermal bremsstrahlung should again provide a good fit to the spectrum. We divided the observations into boxes approximately 3 arcminutes square, resulting in approximately 200–300,000 counts in each box after background subtraction. Modelling the spectrum with thermal bremsstrahlung produces a good fit, with no evidence of any anomalous excess in the 2–2.2 keV region (and, again, the dips in the *Chandra* effective area in this region clearly manifest as a dip in the count rate).

3. Furthermore, if an instrumental or effective area feature at the ten percent level was responsible for this excess, we would expect to also be able to observe it in similar deep observations of other point sources with comparable statistics. To this end we have stacked and analysed the six ACIS-S observations of the quasar 3C273, totalling

180ks of observation time. To make the analysis as similar as possible to our analysis of NGC1275, we have only used observations without a grating, so that the central region of 3C273 is heavily piled up. As with NGC1275, the central pixels are piled-up (even more so than for NGC1275, as there is a readout streak present). Analogously to the case of NGC1275, we have excluded the central region and extracted and analysed the spectra from the wings of the point spread function, with the resulting spectrum having around 90,000 counts after background subtraction. A power-law provides a good fit to the data and no anomalous excess is observed in the region 2–2.2 keV.

These points make it unlikely that the features we observe in the NGC1275 spectrum arise from effective area miscalibrations.

7.3 Gain Miscalibration

Given the high count-rate of NGC1275 and the location of an edge in the effective area around 2 keV, we also check whether the excess we find in the 2–2.2 keV region could be caused by a modification of the gain in the detector response. We investigate this using the `gain` command in `xspec`, which can be used to modify the energy of the response matrix $E' = E/\text{slope} - \text{offset}$, where `slope` and `offset` can be fitted or set to a fixed value. We perform the following checks on the ACIS-I Edge observations:

1. We first examine whether thermal lines are shifted from their correct values. A shift of thermal lines from their known energy values would provide an indication of a need to modify the gain. Since we do not have a high statistics resolution of the 6.4 keV Fe $K\alpha$ line of the AGN spectrum for the ACIS-I data, we instead use thermal lines from the cluster gas emission (which is a background to the AGN spectrum). We extract the thermal spectrum from an annulus region of radii 16 to 19 arcseconds around the AGN. In this spectrum the Fe XXV line at 6.7 keV is visible at very high statistics and at the correct energy, and fitting it with a Gaussian provides no indication of any shift in the gain.
2. We also directly fit the gain using `xspec`. We first fit the absorbed powerlaw together with the gain response parameters `slope` and `offset` over the energy range 0.7–5.0 keV. Compared to fitting purely the absorbed powerlaw, this gives only a marginally improved fit and the best-fit values for `slope` and `offset` are 2-sigma compatible with the standard values 1 and 0. In particular, the excess at 2–2.2 keV remains at the same significance as when we do not fit the gain parameters. As a further check for the possibility that the gain is modified only in the region around 2 keV, we fix the absorbed powerlaw model parameters to their best-fit values from the fit without gain modification and fit only the gain parameters in the region 1.8–2.7 keV. This fit yields best-fit values for the `slope` and `offset` parameters consistent with their standard values within 1-sigma, again leaving the significance of the 2–2.2 keV feature unchanged.

These checks therefore do not provide any suggestion that gain miscalibration accounts for the 2–2.2 keV excess.

7.4 Astrophysical Explanations for 2–2.2 keV excess

Let us now consider potential astrophysical explanations for the excess around 2–2.2 keV.

The first possibility is that the features we observe are not features of the AGN, but instead arises from a mis-subtraction of the background – the vicinity of the AGN is certainly a highly complex environment and there is no reason to expect precise spatial uniformity in the background. However, such an explanation is not viable. The brightness of the AGN, combined with the superb angular resolution of *Chandra*, means that we are able to obtain high ratios of AGN vs. cluster emission. In particular, in both ACIS-I edge and central observations we are able to achieve an AGN to cluster contrast reaching up to 15:1 (and even reaching 60:1 when we model pile-up directly) – making it impossible that a 10% localised excess could arise from the cluster emission.¹⁵

We next consider the possibility that the excess arises from a thermal emission line from ionised gas close to the AGN:

1. Any strong thermal emission present at these energies should also be present as a strong broad excess extending to lower energies. However, for both ACIS-I observations an absorbed power-law is found to be a good fit to the data, implying that any thermal emission has at best a minimal contribution from the spectrum. Even for the ACIS-S observations, where a soft $T \sim 0.9$ keV component is found to be required, the contribution of this component is negligible at $E \sim 2$ keV.
2. In any case, there are no strong emission lines from an ionised gas in this region, either for temperatures $T \sim 0.9$ keV appropriate for the soft excess or for temperatures $T \sim 4$ keV appropriate for the cluster as a whole. The strongest line in the relevant 2–2.2 keV region is a very weak Si XIII line at 2.18 keV (in the rest frame), but this is always much weaker by a factor of ~ 7 than a much stronger Si XIII line at 1.87 keV (in the rest frame). As these are the same ion and the same element, there is no possibility of seeing the weaker line and not the stronger.

We therefore conclude that thermal emission from an ionised gas is not a plausible explanation for the excess emission in the 2–2.2 keV region.

A more substantial consideration – and probably the most promising astrophysical scenario – is fluorescent emission from a neutral $K\alpha$ line. Such neutral $K\alpha$ lines arise from the accretion disc when illuminated with hard X-rays. The Fe $K\alpha$ line at 6.4 keV (in the rest frame) is a well-known feature of the spectrum of an AGN, and is also present in the spectrum of NGC1275. Its presence was reported in the *XMM-Newton* study of [1] as a narrow line at 6.29 keV with a width of 165 eV.

We have checked that the Fe $K\alpha$ line is also present at ~ 6.3 keV in all three cleaned *Chandra* spectra (ACIS-S, ACIS-I edge and ACIS-I midway), with a narrow width consistent with purely instrumental broadening. The best-fit energies for the Fe $K\alpha$ line are 6.29 keV for the cleaned ACIS-S data, 6.33 keV for the cleaned ACIS-I edge data, and 6.35 keV for the cleaned ACIS-I midway observations. Given the uncertainties, these are all mutually

¹⁵ Although the spectra presented for the ACIS-I edge observations involve an AGN to cluster ratio of 6.4:1, by reducing the extraction region we are able to increase this ratio at the cost of only a small reduction in the overall photon count.

consistent and supportive of the fact that the Fe $K\alpha$ line is narrow and redshifted only by the intrinsic Perseus redshift of 0.0176.

While Fe $K\alpha$ emission is the strongest such line, we should consider the possibility of other neutral $K\alpha$ X-ray lines, although as far as we know none have been observed for an unobscured AGN such as NGC1275. The nearest $K\alpha$ lines to the observed feature are from Phosphorus at 1.97 keV (2.01 keV in the rest frame) and Sulphur at 2.27 keV (2.31 keV in the rest frame). The strength of the P line is expected to be negligible [43], but we should consider the S line at 2.27 keV in the detector frame.

An interpretation of the feature at 2–2.2 keV as a S $K\alpha$ line has clear problems. Most obviously, its energy is too high; the excess we observe is consistently at lower energies. In the ACIS-S spectra, with the best energy resolution, no trace of excess emission is seen above 2.2 keV, and for all cases the best-fit central value for the excess (if we naively model it as a Gaussian) is less than 2.2 keV.

This difficulty could be evaded if the S $K\alpha$ energy is modified by large relativistic redshifts associated to the accretion disc. These could arise from a combination of the transverse Doppler effect, relativistic beaming and a general relativistic redshift on climbing out of the black hole potential well. Examples are known where this significantly modifies the form of the Fe $K\alpha$ emission [44], and it is true that the shape of the ACIS-I edge spectra for the 2–2.2 keV excess looks similar to that of the resulting broadened $K\alpha$ lines.

However, when they are present these effects all arise from the motion of the accretion disk around the black hole. As the velocity of particles in Keplerian orbits is independent of their mass, a significant relativistic distortion of the S $K\alpha$ emission should be accompanied by the same distortion of the Fe $K\alpha$ emission. Explaining the 2–2.2 keV feature as S $K\alpha$ emission requires the emission centroid to be shifted by at least 5%. However, such a shift would cause the centroid of the Fe $K\alpha$ line to be located not higher than 6.0 keV – which is incompatible with the observation of the Fe $K\alpha$ line as a narrow line centred at 6.3 keV, with an energy set only by the overall redshift of the Perseus cluster.

There is a second associated difficulty with interpreting the 2–2.2 keV feature as a S $K\alpha$ line: the feature is too strong. From the spectrum in [43] (also in [45] section 19.6), the reflection spectrum from S is expected to be over an order of magnitude smaller than that from Fe (in terms of number of counts). However, the observed excess requires an equivalent number of counts to the Fe line – and so a S $K\alpha$ line far stronger than expectations.

In summary, then, neither thermal ionised emission lines nor fluorescent $K\alpha$ emission appear to provide a satisfactory explanation for the excess in the 2–2.2 keV region.

7.5 Astrophysical Explanations for 3.4–3.5 keV deficit

We now consider possible astrophysical explanations for the feature at 3.4–3.5 keV, present as a deficit of data compared to model. This is present at very high significance in the ACIS-I edge observations, and also at more moderate significance in the ACIS-I midway and ACIS-S observations. In terms of the ratio of data to model, the observations are all compatible.

While this feature lacks the overwhelming statistical significance of the one at 2–2.2 keV, it is harder to come up with instrumental or astrophysical explanations for such a deficit. Compared to the 2–2.2 keV region, at 3.4–3.5 keV the *Chandra* effective areas are smooth functions of energy, and in general this is a clean part of the spectrum.

The simplest explanation is as a statistical fluctuation – although as this deficit is present as a local $\sim 4.5\sigma$ deficit in the ACIS-I edge observation, such an explanation is problematic.

In terms of conventional astrophysical explanations, the most obvious candidate for this deficit would be an atomic absorption line. However it is difficult to see how this could work, as there are no strong lines around this energy. Even more seriously, absorption along the sightline to NGC1275 dominantly arises from the Milky Way – this is evidenced by the fitted values of n_H , which are all consistent with the galactic value towards Perseus of $n_H \simeq 1.5 \times 10^{21} \text{cm}^{-2}$. This implies that, if the feature at 3.4–3.5 keV did arise from atomic absorption, a similar feature should also arise for the continuum cluster spectrum of Perseus, as the galactic n_H values have degree-scale gradients and so n_H is approximately the same for all sightlines to Perseus. However, no such absorption feature at 3.5 keV is detected in the cluster spectra of Perseus — indeed, precisely the opposite is found and instead a significant excess is found at 3.5 keV [46, 47] (for a review see [48]).

While there is consistent support for the 3.4–3.5 keV deficit from the ACIS-S and ACIS-I central observations, by themselves these datasets do not provide any significant statistical evidence. The most plausible conventional explanation of the 3.4–3.5 keV deficit would then be as a localised detector effect for the ACIS-I edge observations, combined with mild downwards statistical fluctuations for the ACIS-S and ACIS-I central observations. However, we also note that inspection of the background region for the ACIS-I edge observations does not show any deficit around 3.4–3.5 keV, and so this would require a rather fortuitous effect at the precise location of NGC1275.

Given the observations of [46, 47] of an unidentified line at 3.5 keV in diffuse emission from the Perseus cluster, it is an interesting fact that we observe a deficit at ~ 3.5 keV in our observations of the NGC1275 AGN at the centre of the cluster. This of course may be just a coincidence, but in the context of new physics models we mention two ways that these facts could be more than simply a coincidence.

The first involves models of excited dark matter invoked for the 3.5 keV line, where dark matter has a resonance at an energy ~ 3.5 keV above its ground state (for example as in [49–54]). In this case there is then an absorption cross section of $E \sim 3.5$ keV photons on dark matter. While for an isotropic initial distribution of photons such absorption and re-emission would not affect the photon spectrum, for a directional beam dark matter absorption will result in an absorption hole in the spectrum (the presence of an absorbing torus around an AGN ensures its outward radiation is indeed directional).

Would such an effect have been observed already elsewhere? We do not see why. In this scenario, the relevant quantity determining the fractional absorption rate is the dark matter column density along the line of sight. It is entirely plausible that the dark matter column density towards the NGC1275 AGN is larger than for almost any other direction in the universe. This is because the emission all originates very close to the central AGN, and so the column density is sensitive to not just the Perseus cluster, but also the central cluster galaxy NGC1275 right down to any sub-pc level dark matter spikes close to the central supermassive black hole. The effect we observe is not large – a 10% reduction over around 100 eV in width – and requires a spectrum with $\mathcal{O}(10^5)$ counts for a statistically significant detection. With a smaller dark matter column density, the effect would reduce to an unobservable $\mathcal{O}(1\%)$ effect.

There is a second possible connection to the 3.5 keV line. An attractive scenario for

the 3.5 keV line involves decay of dark matter to 3.5 keV axion-like particles, which then convert to photons through axion-photon conversion in the cluster magnetic field [55–58]. In this scenario the strength of the 3.5 keV line depends on the efficiency of ALP-photon interconversion – and so broadly is expected to be larger in regions with large magnetic fields extended over wide areas: for example, in galaxy clusters as opposed to galaxies. The 3.5 keV line is observed to be stronger towards the centre of Perseus than for other clusters; one way this could arise is if it fortuitously happens that ALP-photon interconversion is particularly efficient around 3.5 keV for sightlines towards the centre of Perseus. In this case, the presence of a deficit of $E \sim 3.5\text{keV}$ photons from NGC1275 and an excess of $E \sim 3.5\text{keV}$ photons from the cluster as a whole could come from the same underlying physics – efficient photon-ALP interconversion at energies $E \sim 3.5\text{keV}$ along sightlines towards the centre of Perseus.

7.6 ALP interpretation of features

The purpose of this paper was to use the extraordinary dataset of counts from the NGC1275 AGN to search for spectral irregularities, with the intent of constraining ALP parameters. This search has resulted in two features being present in the data at high statistical significance, which we have been unable to account for using conventional instrumental or astrophysical explanations.

These localised oscillatory features are consistent with an ALP origin. If we interpret them as arising from photon-ALP conversion in the cluster magnetic field, then their magnitude allows the ALP-photon coupling $g_{a\gamma\gamma}$ to be determined. While a complete determination would require a precise knowledge of the magnetic field structure along the line of sight from NGC1275 to us, we can estimate approximately that they require an ALP-photon coupling of similar magnitude to the limits placed in the Section 6: $g_{a\gamma\gamma} \sim 1 - 5 \times 10^{-12} \text{ GeV}^{-1}$. For example, Figure 15 shows a fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability $P_{\gamma \rightarrow \gamma}$. In this case, $P_{\gamma \rightarrow \gamma}$ was calculated assuming the existence of ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ and a central magnetic field $B_0 = 25 \mu\text{G}$ (the most optimistic field scenario used in Section 6). We see that the anomalies at 2.2 keV and 3.5 keV have been alleviated by the presence of ALPs (although, for this magnetic field, at the expense of creating similarly sized anomalies at higher energies). Figure 16 shows the corresponding photon survival probability spectrum.

8 Conclusions

The most basic point of this paper is that X-ray observations of NGC1275 are a superb way to search for ALPs. NGC1275 is an extremely bright – and brightening – X-ray point source shining through a galaxy cluster environment. If ALPs exist, they can lead to oscillatory modulations within the energy ranges probed by *Chandra* and *XMM-Newton*. Although there is uncertainty on the precise magnetic field structure along the line of sight to NGC1275, it is almost certainly the case that current and future observations of NGC1275 provide a greater reach in searches for light ALPs with $m_a \lesssim 10^{-12} \text{ eV}$ than even the proposed dedicated experiment IAXO [59].

Previous *Chandra* observations of NGC1275 provide a dataset of extraordinary quality (and, with the potential exception of data taken during *Hitomi*’s pointing at the centre of Perseus, probably the best current dataset for ALP searches). Three factors contribute to




Figure 15. A fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability. We include ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ and assume a central field of $B_0 = 25 \mu\text{G}$. The reduced χ^2 is 1.51, compared to 1.65 for a fit to an absorbed power law without ALPs.




Figure 16. The photon survival probability for the fit shown in Figure 15.

this: firstly, *Chandra*'s angular resolution allows a large contrast between the AGN emission and cluster. Secondly, the existing observations are very deep, and include three independent locations of NGC1275 on the telescope. Thirdly, it a fortuitous fact that in one set of observations, NGC1275 was located on the edge of the chip, thereby providing a clean observation which minimises pile-up.

Nonetheless, from the perspective of ALP physics the dataset could be substantially improved with even relatively modest further observations. All existing observations were performed with the nominal frame time of ~ 3 seconds, but it is possible to reduce the frame

time as low as 0.2 seconds, which would reduce pile-up by an order of magnitude. Given that NGC1275 is brighter now than it was in 2009, further off-axis observations of NGC1275 with reduced frame time and exposures of $\sim 100 - 200$ ks would result in both a larger and cleaner dataset than the best used in this paper, the 2009 ACIS-I edge observations.

In this paper we have used all existing observations of NGC1275 with *Chandra* and *XMM-Newton* to search for spectral modulations induced by ALPs. If they exist, ALPs give rise to spectral modulations, and we have used the absence of modulations at the $\mathcal{O}(30\%)$ level to place leading current bounds on ALP-photon couplings. At the $\mathcal{O}(10\%)$ level, there are two main modulations in the data – one upward around 2–2.2 keV, and one downward around 3.4–3.5 keV. The former is observed in all *Chandra* datasets, and is present with overall statistical significance well in excess of 5σ . Pile-up is a concern for such a bright source, but all measures we have taken to reduce pile-up leave this modulation present and at the same strength (the statistical significance of course decreases as the amount of data used is reduced). The downward modulation at 3.4–3.5 keV is not as large, but still has 4.5σ local statistical significance in the full ACIS-I edge dataset. We have been unable to identify conventional instrumental or astrophysical processes that can give rise to either of these effects.

XMM-Newton observations are consistent with these results, but do not provide additional supporting evidence. *XMM-Newton* is a complementary instrument to *Chandra*, and for this purpose it is not as sensitive. In particular, the advantages of *XMM-Newton*’s high effective area compared to *Chandra* are outweighed by its worse angular resolution, making it harder to distinguish the AGN emission from the cluster emission. As all *XMM-Newton* observations are on-axis, pile-up is also a significant contaminant for these observations.

In summary, axion-like particles are one of the most plausible ways to extend the Standard Model, and X-ray observations of the centre of the Perseus cluster provide an outstanding way to look for them. In this respect *Chandra* is not only an X-ray astronomy satellite of the highest quality, but possibly also a discovery tool for fundamental physics.

Acknowledgments

We thank David Marsh for comments on the draft, Alexis Finoguenov and Jeroen Franse for helpful discussions, and the Bethe Centre for Theoretical Physics for hospitality while part of this work was carried out. We also thank the staff and science teams of the *Chandra* and *XMM-Newton* X-ray observatories for wonderful instruments, abundant documentation that is sufficiently clear that particle theorists can follow it, and friendly helpdesks. This project is funded in part by the European Research Council starting grant ‘Supersymmetry Breaking in String Theory’ (307605). JC is also funded by a Royal Society University Research Fellowship.

A Ultraclean Spectra

Here we also show spectra extracted for the ACIS-I edge observations using an ultra-clean data sample. The purpose of this is to provide further evidence that the spectral features observed at 2 - 2.2 keV and 3.4 - 3.5 keV are present even in the cleanest possible data samples available. For this purpose, we have extracted the data from the region shown in Figure 17. This uses the data from the the ACIS-I edge observations where the source is




Figure 17. The extraction region used to provide the cleanest possible data sample. The red strikethrough line corresponds to an excluded central ellipse.

highly off-axis, and furthermore extracts only from deep within the wings of the point spread function. After background subtraction there are 36000 counts in the spectrum (47000 prior to background subtraction); this involves a cut of more than 80% of the counts from the AGN compared to the full ACIS-I edge extraction.

The spectrum is shown in Figure 18. The background-subtracted data is fit to an absorbed powerlaw between 0.8 and 5 keV, grouping counts to a minimum of 350 per bin. The resulting fit produces $n_H = 1.53 \times 10^{21} \text{cm}^{-2}$ and an index of $\gamma = 1.69$. The features seen in the full data are still present (local excesses of 1.9 and 3.1 sigma around 2.2 keV, and a local deficit of 2.2 sigma at 3.4-3.5 keV), and the magnitude of these features is entirely compatible with the results found in the larger data samples.

References

- [1] E. Churazov, W. Forman, C. Jones, and H. Bohringer, *XMM-newton observations of the perseus cluster I: the temperature and surface brightness structure*, *Astrophys. J.* **590** (2003) 225–237, [[astro-ph/0301482](#)].
- [2] S. Yamazaki, Y. Fukazawa, M. Sasada, R. Itoh, S. Nishino, H. Takahashi, K. Takaki, K. S. Kawabata, M. Yoshida, and M. Uemura, *X-Ray and Optical Monitoring of a Gamma-Ray-Emitting Radio Galaxy, NGC 1275*, *Publications of the Astronomical Society of Japan* **65** (Apr., 2013).
- [3] B. Balmaverde, A. Capetti, and P. Grandi, *The Chandra view of the 3c/fri sample of low luminosity radio-galaxies*, *Astron. Astrophys.* **451** (2006) 35, [[astro-ph/0601175](#)].
- [4] A. C. Fabian, S. A. Walker, C. Pinto, H. R. Russell, and A. C. Edge, *Effects of the variability of the nucleus of NGC 1275 on X-ray observations of the surrounding intracluster medium*, *Mon. Not. Roy. Astron. Soc.* **451** (2015), no. 3 3061–3067, [[1505.03754](#)].
- [5] J. P. Conlon, *The QCD axion and moduli stabilisation*, *JHEP* **05** (2006) 078, [[hep-th/0602233](#)].
- [6] P. Svrcek and E. Witten, *Axions In String Theory*, *JHEP* **06** (2006) 051, [[hep-th/0605206](#)].




Figure 18. ACIS-I observations with an ultra-clean extraction region, fit to an absorbed power law between 0.8 and 5 keV.

- [7] M. Cicoli, M. Goodsell, and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *JHEP* **10** (2012) 146, [[1206.0819](#)].
- [8] A. Ringwald, *Exploring the Role of Axions and Other WISPs in the Dark Universe*, *Phys. Dark Univ.* **1** (2012) 116–135, [[1210.5081](#)].
- [9] J. W. Brockway, E. D. Carlson, and G. G. Raffelt, *SN1987A gamma-ray limits on the conversion of pseudoscalars*, *Phys. Lett.* **B383** (1996) 439–443, [[astro-ph/9605197](#)].
- [10] J. A. Grifols, E. Masso, and R. Toldra, *Gamma-rays from SN1987A due to pseudoscalar conversion*, *Phys. Rev. Lett.* **77** (1996) 2372–2375, [[astro-ph/9606028](#)].
- [11] A. Payez, C. Evoli, T. Fischer, M. Giannotti, A. Mirizzi, and A. Ringwald, *Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles*, *JCAP* **1502** (2015), no. 02 006, [[1410.3747](#)].
- [12] P. Sikivie, *Experimental Tests of the Invisible Axion*, *Phys. Rev. Lett.* **51** (1983) 1415–1417. [Erratum: *Phys. Rev. Lett.* 52,695(1984)].
- [13] G. Raffelt and L. Stodolsky, *Mixing of the Photon with Low Mass Particles*, *Phys. Rev.* **D37** (1988) 1237.
- [14] C. Burrage, A.-C. Davis, and D. J. Shaw, *Active Galactic Nuclei Shed Light on Axion-like-Particles*, *Phys. Rev. Lett.* **102** (2009) 201101, [[0902.2320](#)].
- [15] J. P. Conlon and M. C. D. Marsh, *Excess Astrophysical Photons from a 0.1-1 keV Cosmic Axion Background*, *Phys. Rev. Lett.* **111** (2013), no. 15 151301, [[1305.3603](#)].
- [16] D. Wouters and P. Brun, *Constraints on Axion-like Particles from X-Ray Observations of the Hydra Galaxy Cluster*, *Astrophys. J.* **772** (2013) 44, [[1304.0989](#)].
- [17] S. Angus, J. P. Conlon, M. C. D. Marsh, A. J. Powell, and L. T. Witkowski, *Soft X-ray Excess in the Coma Cluster from a Cosmic Axion Background*, *JCAP* **1409** (2014), no. 09 026, [[1312.3947](#)].

- [18] J. P. Conlon, M. C. D. Marsh, and A. J. Powell, *Galaxy Cluster Thermal X-Ray Spectra Constrain Axion-Like Particles*, [1509.06748](#).
- [19] D. Wouters and P. Brun, *Irregularity in gamma ray source spectra as a signature of axionlike particles*, *Phys. Rev.* **D86** (2012) 043005, [[1205.6428](#)].
- [20] P. Brax, P. Brun, and D. Wouters, *Galaxy cluster constraints on the coupling to photons of low-mass scalars*, *Phys. Rev.* **D92** (2015) 083501, [[1505.01020](#)].
- [21] **The Fermi-LAT Collaboration**, *Search for spectral irregularities due to photon-axion-like particle oscillations with the Fermi Large Area Telescope*, [1603.06978](#).
- [22] P. Sikivie, *Detection rates for “invisible”-axion searches*, *Phys. Rev. D* **32** (Dec, 1985) 2988–2991.
- [23] J. W. Brockway, E. D. Carlson, and G. G. Raffelt, *SN1987A gamma-ray limits on the conversion of pseudoscalars*, *Phys. Lett.* **B383** (1996) 439–443, [[astro-ph/9605197](#)].
- [24] J. A. Grifols, E. Masso, and R. Toldra, *Gamma-rays from SN1987A due to pseudoscalar conversion*, *Phys. Rev. Lett.* **77** (1996) 2372–2375, [[astro-ph/9606028](#)].
- [25] M. Fairbairn, T. Rashba, and S. V. Troitsky, *Photon-axion mixing and ultra-high-energy cosmic rays from BL Lac type objects - Shining light through the Universe*, *Phys. Rev.* **D84** (2011) 125019, [[0901.4085](#)].
- [26] A. Mirizzi and D. Montanino, *Stochastic conversions of TeV photons into axion-like particles in extragalactic magnetic fields*, *JCAP* **0912** (2009) 004, [[0911.0015](#)].
- [27] F. Tavecchio, M. Roncadelli, G. Galanti, and G. Bonnoli, *Evidence for an axion-like particle from PKS 1222+216?*, *Phys. Rev.* **D86** (2012) 085036, [[1202.6529](#)].
- [28] A. Payez, J. R. Cudell, and D. Hutsemekers, *New polarimetric constraints on axion-like particles*, *JCAP* **1207** (2012) 041, [[1204.6187](#)].
- [29] D. Horns, L. Maccione, M. Meyer, A. Mirizzi, D. Montanino, and M. Roncadelli, *Hardening of TeV gamma spectrum of AGNs in galaxy clusters by conversions of photons into axion-like particles*, *Phys. Rev.* **D86** (2012) 075024, [[1207.0776](#)].
- [30] M. Meyer, D. Horns, and M. Raue, *First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations*, *Phys. Rev.* **D87** (2013), no. 3 035027, [[1302.1208](#)].
- [31] A. J. Powell, *A Cosmic ALP Background and the Cluster Soft X-ray Excess in A665, A2199 and A2255*, *JCAP* **1509** (2015), no. 09 017, [[1411.4172](#)].
- [32] A. Dobrynina, A. Kartavtsev, and G. Raffelt, *Photon-photon dispersion of TeV gamma rays and its role for photon-ALP conversion*, *Phys. Rev.* **D91** (2015) 083003, [[1412.4777](#)].
- [33] M. Schleder and G. Sigl, *Constraining ALP-photon coupling using galaxy clusters*, *JCAP* **1601** (2016), no. 01 038, [[1507.02855](#)].
- [34] A. Bonafede, L. Feretti, M. Murgia, F. Govoni, G. Giovannini, D. Dallacasa, K. Dolag, and G. B. Taylor, *The Coma cluster magnetic field from Faraday rotation measures*, *Astron. Astrophys.* **513** (2010) A30, [[1002.0594](#)].
- [35] G. B. Taylor, N. E. Gugliucci, A. C. Fabian, J. S. Sanders, G. Gentile, and S. W. Allen, *Magnetic fields in the center of the perseus cluster*, *Mon. Not. Roy. Astron. Soc.* **368** (2006) 1500–1506, [[astro-ph/0602622](#)].
- [36] A. Fruscione, J. C. McDowell, G. E. Allen, N. S. Brickhouse, D. J. Burke, J. E. Davis, N. Durham, M. Elvis, E. C. Galle, D. E. Harris, D. P. Huenemoerder, J. C. Houck, B. Ishibashi, M. Karovska, F. Nicastro, M. S. Noble, M. A. Nowak, F. A. Primini, A. Siemiginowska, R. K. Smith, and M. Wise, *CIAO: Chandra’s data analysis system*, in

- Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 6270 of *Proceedings of the International Society for Optical Engineering*, p. 62701V, June, 2006.
- [37] P. Freeman, S. Doe, and A. Siemiginowska, *Sherpa: a mission-independent data analysis application*, in *Astronomical Data Analysis* (J.-L. Starck and F. D. Murtagh, eds.), vol. 4477 of *Proceedings of the International Society for Optical Engineering*, pp. 76–87, Nov., 2001. [astro-ph/0108426](#).
 - [38] J. E. Davis, *Event Pileup in Charge-coupled Devices*, *Astrophys. J* **562** (Nov., 2001) 575–582.
 - [39] K. A. Arnaud, *XSPEC: The First Ten Years*, in *Astronomical Data Analysis Software and Systems V* (G. H. Jacoby and J. Barnes, eds.), vol. 101 of *Astronomical Society of the Pacific Conference Series*, p. 17, 1996.
 - [40] R. Morrison and D. McCammon, *Interstellar photoelectric absorption cross-sections, 0.03–10 keV*, *Astrophys. J.* **270** (1983) 119.
 - [41] ESA: XMM-Newton SOC, *XMM-Newton Users Handbook, Issue 2.13.1* (2016).
 - [42] V. Vacca, M. Murgia, F. Govoni, L. Feretti, G. Giovannini, R. A. Perley, and G. B. Taylor, *The intracluster magnetic field power spectrum in A2199*, *Astron. Astrophys.* **540** (2012) A38, [[1201.4119](#)].
 - [43] A. Fabian, *Emission Lines: Signatures of Relativistic Rotation*, in *Theory of Black Hole Accretion Disks*, ed. Abramowicz, M. A. (CUP) (1998).
 - [44] Y. Tanaka, K. Nandra, A. C. Fabian, H. Inoue, C. Otani, T. Dotani, K. Hayashida, K. Iwasawa, T. Kii, H. Kunieda, F. Makino, and M. Matsuoka, *Gravitationally redshifted emission implying an accretion disk and massive black hole in the active galaxy MCG-6-30-15*, *Nature* **375** (June, 1995) 659–661.
 - [45] M. S. Longair, *High Energy Astrophysics*. CUP, 2011.
 - [46] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, *Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters*, *Astrophys. J.* **789** (2014) 13, [[1402.2301](#)].
 - [47] A. Boyarsky, O. Ruchayskiy, D. Iakubovskiy, and J. Franse, *Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster*, *Phys. Rev. Lett.* **113** (2014) 251301, [[1402.4119](#)].
 - [48] D. Iakubovskiy, *Observation of the new line at 3.55 keV in X-ray spectra of galaxies and galaxy clusters*, [[1510.00358](#)].
 - [49] D. P. Finkbeiner and N. Weiner, *An X-Ray Line from eXciting Dark Matter*, [[1402.6671](#)].
 - [50] J. M. Cline and A. R. Frey, *Nonabelian dark matter models for 3.5 keV X-rays*, *JCAP* **1410** (2014) 013, [[1408.0233](#)].
 - [51] J. M. Cline and A. R. Frey, *Consistency of dark matter interpretations of the 3.5 keV x-ray line*, *Phys. Rev.* **D90** (2014), no. 12 123537, [[1410.7766](#)].
 - [52] A. Berlin, A. DiFranzo, and D. Hooper, *3.55 keV line from exciting dark matter without a hidden sector*, *Phys. Rev.* **D91** (2015), no. 7 075018, [[1501.03496](#)].
 - [53] P. Agrawal, Z. Chacko, C. Kilic, and C. B. Verhaaren, *A Couplet from Flavored Dark Matter*, *JHEP* **08** (2015) 072, [[1503.03057](#)].
 - [54] F. D’Eramo, K. Hambleton, S. Profumo, and T. Stefaniak, *Dark Matter Inelastic Up-Scattering with the Interstellar Plasma: An Exciting New Source of X-Ray Lines, including at 3.5 keV*, [[1603.04859](#)].
 - [55] M. Cicoli, J. P. Conlon, M. C. D. Marsh, and M. Rummel, *3.55 keV photon line and its morphology from a 3.55 keV axionlike particle line*, *Phys. Rev.* **D90** (2014) 023540, [[1403.2370](#)].

- [56] J. P. Conlon and F. V. Day, *3.55 keV photon lines from axion to photon conversion in the Milky Way and M31*, *JCAP* **1411** (2014) 033, [[1404.7741](#)].
- [57] J. P. Conlon and A. J. Powell, *A 3.55 keV line from $DM \rightarrow a \rightarrow \gamma$: predictions for cool-core and non-cool-core clusters*, *JCAP* **1501** (2015), no. 01 019, [[1406.5518](#)].
- [58] P. D. Alvarez, J. P. Conlon, F. V. Day, M. C. D. Marsh, and M. Rummel, *Observational consistency and future predictions for a 3.5 keV ALP to photon line*, *JCAP* **1504** (2015), no. 04 013, [[1410.1867](#)].
- [59] E. Armengaud *et. al.*, *Conceptual Design of the International Axion Observatory (IAXO)*, *JINST* **9** (2014) T05002, [[1401.3233](#)].